A USDOT NATIONAL UNIVERSITY TRANSPORTATION CENTER

Creating and Integrating Solutions to Enable the 'Complete Trip'

Stephen F. Smith (PI), https://orcid.org/0000-0002-7053-3166

Patrick Carrington, https://orcid.org/0000-0001-8923-0803
Artur Dubrawski, https://orcid.org/0000-0002-2372-0831
Srinivas Narasimhan, https://orcid.org/0000-0003-0389-1921
Jean Oh, https://orcid.org/0000-0001-9709-2658
Zachary B. Rubinstein, https://orcid.org/0000-0002-6344-8692
Robert Tamburo, https://orcid.org/0000-0002-5636-9443
Ji Zhang, https://orcid.org/0000-0002-4692-5645

FINAL RESEARCH REPORT

Contract # 69A3551747111

Creating and Integrating Solutions to Enable the 'Complete Trip'

Final Report

Stephen F. Smith[1] (PI),
Patrick Carrington[2], Artur Dubrawski[1], Shrinivas Narasimhan[1],
Jean Oh[1], Zachary B. Rubinstein[1], Robert Tamboro[1], Ji Zhang[1]

[1] The Robotics Institute [2] Human Computer Interaction Institute

January 14, 2023

1. Overview

In this report, we summarize the progress and accomplishments achieved in Mobility21 University Transportation Center Award #40459-8.3-1080376, titled "Creating and Integrating Solutions to Enable the 'Complete Trip'". One integrating theme that has been promoted in recent years for framing the technology needs of mobility-challenged individuals is that of facilitating the "Complete Trip" [NCMM 2020]. Toward this end, this project has focused on two chief deliverables: (1) a mobile smartphone app for persons with disabilities that integrates pedestrian-friendly route planning, real-time navigation, safe intersection crossing, and coordination of multi-modal travel legs through cloud-based traveler-to-infrastructure (T2I) communication, and (2) the development of enabling technology for an intelligent wheelchair that couples dynamic obstacle detection and autonomous navigation capabilities to provide wheelchair users with real-time driver assist for traversing challenging terrain (e.g., intersection curb cuts, potholes, sidewalks under repair). The project focused specifically on complete trip support for pedestrians with two types of disabilities, wheelchair users and vision-impaired individuals, but we believe the capabilities that have been developed will also be applicable and useful to other mobility-challenged individuals.

Toward the first objective of developing a smartphone app for complete trip support, the project has accomplished the following:

- In collaboration with industrial partners pathVu and Rapid Flow Technologies, we have created ***PedPal+***, a smartphone app that combines and extends two previously developed component technologies:
    - *PedPal* – A smartphone app designed to exploit real-time communication with the traffic signal system at the intersection to promote safe and efficient crossing of signalized intersections [Smith et.al 2019].
    - *pathVu Mobility Planner* – A smartphone app that takes advantage of an underlying sidewalk mapping data base to generate the most accessible

route from A to B (as opposed to the fastest) and then provides real-time navigation support as the route is executed [pathVu 2022].

PedPal+ is designed to allow its user to request the most accessible route from one location to another (e.g., from home to a bus stop) at the beginning of a travel episode, and then tracks the user and gives navigation cues as the user begins to execute. As the user approaches a signalized intersection along the route, the app switches to safe intersection crossing mode, and begins interacting with the traffic signal system to gather information about when different crossing options are scheduled to occur. The app then indicates to its user when the desired crossing direction will next be active (green), it communicates to the traffic signal system how much time its user needs to make the cross (using intersection geometry information and knowledge of the user's speed) and prepares the user for crossing as the scheduled time approaches.

- To provide a reliable basis for detecting when the user arrives at a crossing corner, a straightforward geofencing solution was developed based on the placement of an *Ultra-Wide Band (UWB)* beacon at each corner light pole of the intersection. By establishing communication between the *PedPal* app and these beacons, range data is used to both recognize corners and provide accurate tracking of progress during street crossing. Overall, the introduction of UWB beacons as additional infrastructure provides 2-3 cm localization accuracy within the app and solves the longstanding smartphone localization problem that has limited *PedPal* capabilities to date.

- Further exploiting the presence of UWB beacons at each corner of the intersection, triangularization algorithms for detecting pedestrian movement outside of the crosswalk during crossing and issuing corrective course alerts were developed and validated using the SUMO traffic simulator. Intended principally for vision impaired pedestrians, the concept is to issue a haptic alert (e.g., 1 vibration for a left correct and 2 for a right correct).

- Finally, the ***PedPal+*** app was extended to utilize real-time T2I communication to broadcast crosswalk presence information to all approaching "connected" vehicles to increase their awareness of the pedestrian crossing. In collaboration with *Argo AI*, control of their autonomous vehicles was extended to receive these messages and take precautionary action (i.e., slow down) if appropriate to provide a demonstration of this capability.

- All of the above capabilities were consolidated within the app and incorporated into a live demonstration with a vision-impaired pedestrian traveling from his residence at Spirit Avenue and Highland Avenue in East Liberty, through the intersection of Highland Avenue and Centre Avenue and onto the east-bound bus stop at the corner of Penn Avenue and Highland Avenue.

Toward the second objective of an intelligent autonomous wheelchair, the project has produced the following additional accomplishments:

- A full-stack autonomous wheelchair was developed to provide a platform for investigating the development of core techniques for outdoor navigation and dynamic obstacle avoidance.

- A Self-supervised Learning approach to *Traversability analysis* was developed and applied to generate a cost map based on visual features of different types of traversable terrain. The model was learned using first person-view camera images and 3D lidar scans to map the first-person view into traversability by using a robot's internal vehicle state as a signal for the ride comfort. This result was then integrated into an autonomous wheelchair navigation algorithm to provide a basis for outdoor wheelchair driving. To provide a baseline for evaluation, a second semantic classification approach was also developed and used to quantify the leverage provided by the traversability-based approach. A user study showed the proposed approach to be superior to the baseline with respect to the safety, stability, trustworthiness, and overall preference.

- To enable autonomous wheelchair navigation along routes that require traversal of both indoor and outdoor spaces, the above traversability-based approach was subsequently combined with a more traditional indoor autonomous navigation approach based on object recognition and scene analysis. This broader capability was validated by demonstrating the ability of a prototype intelligent wheelchair to travel autonomously from the 3rd floor Atrium inside Newell and Simon Hall on the CMU Campus, out onto the sidewalks running between and around other CMU buildings, and eventually arriving at a sidewalk nearby a former bus stop on Forbes Avenue.

- Finally, light curtain technology was developed for dynamic detection of obstacles during autonomous navigation. Specifically, techniques were developed that uses a light curtain (1) to track surrounding objects and estimate their velocity, based on particle filtering, and (2) to estimate the sensor ego-motion, by associating range data to visual features and tracking the visual features. Both techniques were tested and validated on board the prototype intelligent wheelchair.

In the remainder of this report, we describe accomplishments in all areas in more detail and discuss next steps for technology maturation and transition.

2. ***PedPal+***: A Smartphone App for Support of the Complete Trip

The ***PedPal+*** app developed over the course of this project consolidates and extends previously developed component technologies for safe intersection crossing (the original PedPal smartphone app [Smith et.al 2019]) and pedestrian-friendly routing (pathVu's current wheelchair user routing and navigation app [pathVu 2022]). Methodologically, we have taken the PedPal smartphone app as our starting point. PedPal derives its core functionality from its real-time connectivity to Surtrac [Xie et.al, 2012, Smith et.al 2013, Smith2020], a decentralized adaptive traffic signal control system developed originally at CMU and now provided commercially by Rapid Flow Technologies. PedPal relies on interaction with Surtrac to obtain intersection geometry and signal phase information, to ensure that the pedestrian is given enough time to cross when the crossing phase becomes active, and to dynamically extend the crossing time if the pedestrian is observed to be moving slower than expected. Over the course of this project, we have improved and extended its safe intersection crossing capabilities, and then expanded its scope by importing routing capabilities from pathVu's routing app.

In the following subsections, we first highlight the major technical advances that have been incorporated into PedPal to create the current ***PedPal+*** app. We then present a representative use case to illustrate how the app is intended to be used. Finally, we describe some current limitations and next steps for rounding out and hardening the ***PedPal+*** technology, as well as strategies for subsequent transition of the technology into practice.

2.1 Boosting App Localization Capabilities through Ultra-Wide Band Sensors

One of the big challenges in the development of the original PedPal app was achieving sufficient localization accuracy to provide the necessary pedestrian tracking capabilities to support vision-impaired individuals. The iPhone was chosen as the implementation platform because of the strength of its native accessibility features (voice-over, zoom, etc.). However, like all contemporary smartphone technology, the iPhone's native localization algorithm quickly proved too approximate for two basic needs: (1) identifying pedestrian arrival at an intersection corner and (2) real-time tracking of progress during pedestrian crossing (to determine whether a real-time extension of the current green phase should be requested). To address this problem, the original PedPal development effort turned to additional infrastructure support. Bluetooth beacons were introduced on the pole at each corner of the intersection and configured to broadcast range information to the smartphone app. A geofence is then defined around the known location of each corner beacon (which is accessible as extra information added to the intersection's MAP message, which provides the intersection geometry). In practice, we were able to get repeatable corner detection using a 10m radius geofence (which still leaves vision-impaired individuals with considerable uncertainty).

With the emergence of Apple's Air Tag tracking sensors in early 2022, the opportunity to exploit ultra-wide band (UWB) to provide a better solution to corner detection became apparent. To explore this possibility, we acquired a set of iPhone 12 mini smartphones (the first iPhone to embed UWB radio capability) for use as experimental UWB "beacons", and utilized Apple's *nearby interaction* framework to establish connectivity between all corner beacons and the *PedPal*+ app. In field test experiments at the intersection of Highland Avenue and Centre Avenue in East Liberty, we have verified the advertised 2-3cm localization accuracy in ideal conditions, allowing us to reduce the corner geofence to a 3m radius with reliable results. This capability also provides a basis for accurately tracking pedestrian progress and strengthens PedPal's original ability to dynamically extend the crossing phase if the pedestrian is moving slower than expected. More recently, Estimote has begun producing an inexpensive UWB beacon (approximate cost: $30 each) that utilizes the same *nearby interaction* communication framework, and one ongoing thread of work focuses on field testing this sensing alternative.

2.2 Detecting Pedestrian Movement outside of the Crosswalk

Figure 1: Estimating pedestrian location during crossing.

The localization accuracy provided by UWB technology has also introduced the opportunity to introduce new safe intersection crossing capabilities. Specifically, we have developed a *pedestrian localization and reorientation algorithm* capable of detecting pedestrian movement outside of the crosswalk and signaling a corrective alert when this happens. [Hata et.al 2022]. Starting from the assumption that the orientation of the pedestrian at the origin corner prior to crossing is known, the algorithm repeatedly triangulates the range information communicated by the three closest intersection beacons to estimate the pedestrian's location at discrete time steps (e.g., every second) once crossing commences (see Figure 1). At each time step, the last *n* locations in this time series data are used by the algorithm to predict the pedestrian's future trajectory. If this predicted trajectory crosses the crosswalk boundary (or some specified minimum distance from the crosswalk boundary), then a haptic alert is issued by the app – one vibration for "move left to correct" and two for "move right" – to reorient the pedestrian, and the predicted trajectory is updated accordingly.

To evaluate both localization and reorienting aspects of the algorithm, A SUMO pedestrian traffic simulation model was developed with a simple (linear) noise model, where the amount of noise that is added to the actual range information communicated by a given beacon increases (or decreases) as the pedestrian gets farther (or closer) to the beacon. In a series of experiments with a model of the intersection of Centre Avenue and Highland Avenue in the East Liberty region of Pittsburgh PA and a walking model for pedestrians, the localization method was found to provide high accuracy, differing by only 0.003% from the ground truth, and the reorienting algorithm was shown to consistently result in successful crossing within the crosswalk. Example trajectories and recoveries generated for west to east crossings of the intersection are shown in Figure 2.

(a) (b)

Figure 2: Recognition and Correction of pedestrian drifting (a) south, (b) north

We are currently working toward validating this capability in the field at the actual intersection of Centre Avenue and Highland Avenue.

2.3 Signaling Crosswalk Presence to Approaching Connected Vehicles

A separate extension to the original PedPal app that was developed over the course of the project involved providing the app with the ability to signal the expected presence of its user in the crosswalk to any approaching connected vehicle. This real-time *“PedPal+ to Vehicle”* communication capability was provided by re-engineering and expanding the PedPal cloud server originally developed to provide “*PedPal to Infrastructure*” connectivity with the Surtrac traffic signal control system.

The expanded PedPal cloud server protocols are illustrated graphically in Figure 3. As the *PedPal*+ app becomes within range of an intersection, it receives the standard MAP message from the cloud server, followed by a continuous stream of Signal Phase and Timing (SPaT) messages that are generated by the Surtrac process running at the intersection. The MAP message provides information on the intersection geometry (width of lanes, number of lanes, crosswalk IDs/locations, corner locations, etc.). Each SPaT message provides updated information about the intersection’s signal phases (e.g., currently active crossing phase, time until future crossing phases, etc.). When the pedestrian arrives at an intersection corner and selects the crossing direction, PedPal+ communicates a Signal Request Message (SRM) through the cloud server to the Surtrac process that specifies the amount of time needed for the pedestrian to cross when the desired crossing phase becomes active (given knowledge of the pedestrian’s travel speed and the geometry of the intersection). After responding to this request, Surtrac sends a Signal Status Message (SSM) back through the server to the app, indicating whether the crossing time was extended (i.e., that it was not the case that the requested extension results in violation of the maximum time constraint associated with the current active phase).

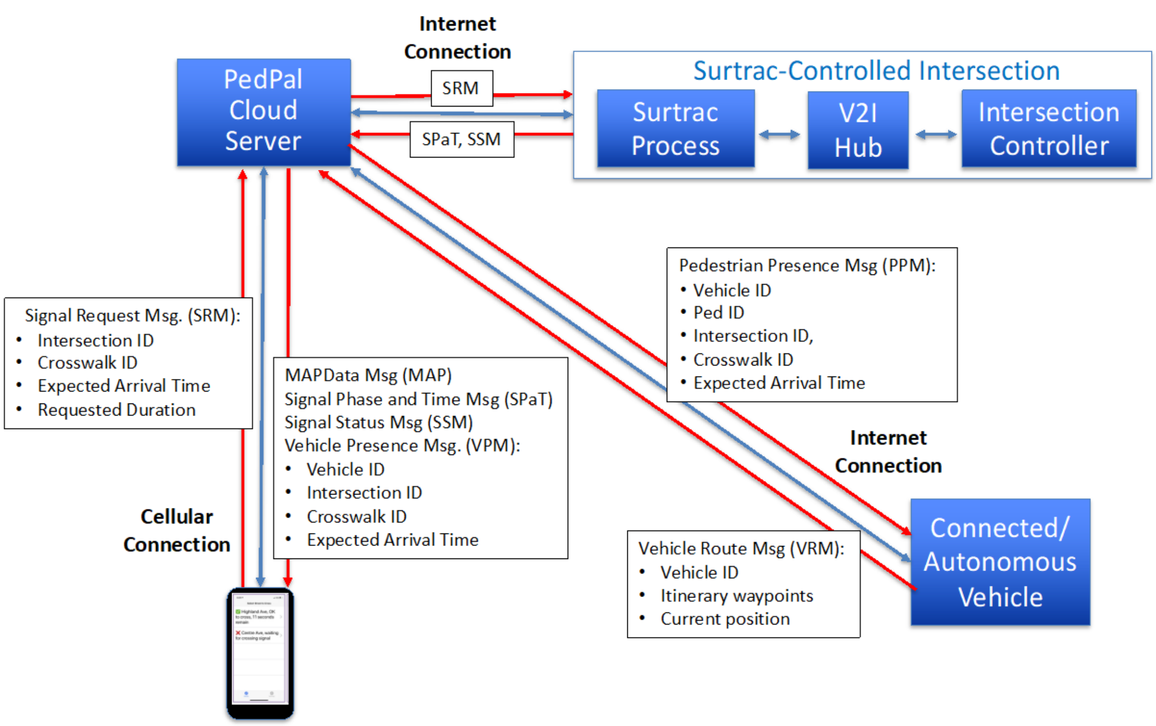

Figure 3: V-to-I-to-*PedPal*+ Connectivity

To make *PedPal+* aware of vehicles in the area that have real-time connectivity, we adopt the assumption that connected vehicles are incentivized to share their routes with the infrastructure and define a Vehicle Route Message (VRM). Rationale for this assumption stems from previous research, which has shown that if a vehicle is willing to share their routes with a Surtrac controlled traffic network, this information can be used to move this vehicle through the network substantially faster than it would without sharing this information [Hawkes16, Smith21].[1] As a connected vehicle enters a Surtrac controlled network, it communicates a VRM to the *PedPal+* cloud server, providing a vehicle ID, a sequence of waypoints characterizing its planned route through the network and the vehicle's current location. The server then uses the sequence of waypoints, together with travel speed information, to estimate the arrival times of the vehicle at each successive signalized intersection along its planned route. These predicted arrival times are stored locally for later use in determining the proximity of various connected vehicles to a given intersection at a particular future point in time.

When the *PedPal+* app informs a pedestrian waiting at an intersection corner that the crossing phase is now active and it is ok to cross, the app also communicates a newly defined Pedestrian Presence Message (PPM) to the cloud server, indicating the Intersection ID, Crosswalk ID, and Expected Arrival Time at the destination corner. Upon receipt of the PPM, the cloud server collects the set of connected vehicles whose predicted arrival times at the designated intersection are in proximity to the crossing pedestrian and relays the PPM to each of these vehicles.

To demonstrate this capability, an API was developed to allow message passing between the ***PedPal+*** cloud server and the cloud server used by Argo AI to maintain the routes and real-time status of their autonomous vehicle fleet as they drive around Pittsburgh neighborhoods. The above summarized handling of VRM and PPM message types was implemented on both sides of this API and experimental code for responding to the receipt of a PPM onboard a vehicle was created. This code was designed to command the vehicle to slow down and stop if necessary to avoid entry into the crosswalk during the specified pedestrian presence interval. For our testing purposes, a specific autonomous vehicle running this experimental code was then dispatched to drive a circular route that passed through our test intersection of Centre Avenue and Highland Avenue (see Figure 5 below). The route called for a right turn onto Centre from Highland. Once the vehicle was on site and driving this route, several pedestrian crosses of Centre Avenue were carried out, each starting from the southeast corner of the intersection as the Argo AI vehicle was either approaching the intersection or already stopped at the intersection waiting for the green. For those cases where there

[1] This reduction in travel time is on top of the benefit accrued from adopting Surtrac real-time adaptive control.

was temporal overlap in the presence interval and the vehicle's right turn, it was verified by the person riding in the vehicle that the vehicle speed was reduced to eliminate the possibility of a collision.

2.4 Integrating Accessible Routing

Moving beyond the scope of the safe intersection crossing application, a final capability that was added to the ***PedPal+*** app is the ability to generate and navigate *accessible* pedestrian routes. This extension was built directly on technology developed originally by pathVu for its wheelchair navigation app. This technology, at its core, is a cloud service for generating accessible routes. An accessible route is defined as one that promotes safe, unobstructed passage from point a to point b, in contrast to routing systems such as Google Maps, Waze or Apple Maps, which generate routes based on travel time objectives. To support the generation of "pedestrian friendly" routes, the pathVu generator relies on a continually maintained *sidewalk database* that is constructed by visually mapping neighborhood sidewalks for cracked, crumbling or no-existent pavement, and then crowdsourcing transient obstructions such as pavement under construction or adjacent building work zones. The route generation algorithm combines this information with traffic volume and intersection geometry characteristics to evaluate alternative routes from origin point a to destination point b and select the most accessible.

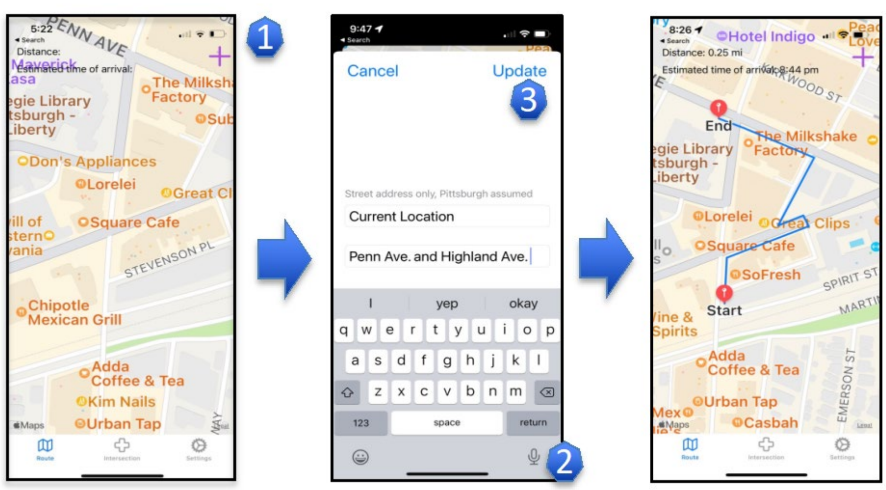

(a) (b) (c)

Figure 4: Requesting an Accessible Route: (1) Select "+" from initial route view screen to trigger text dialog screen, (2) select microphone to dictate destination location, (3) select update to trigger route generator and display the result.

To exploit this service, an API was developed and implemented to interface **PedPal+** with pathVu's route generation cloud service. An additional *route* mode was added to PedPal's original *intersection* and *settings* modes (see Figure 4.a below). From this base interface screen, it is possible to request a new accessible route. This brings up a text screen for expressing the desired origin and destination of the route (see Figure 4.b). By default, the origin location is assumed to be the current location of the app. If the app user is vision impaired, Apple IOS's native dictation capability is used to specify the intended destination. Once specified the remote route generator is invoked, and the returned route is displayed via Apple Maps (Figure 4.c). Along with the route, the generator also provides the computed route distance and the expected arrival time at the destination if the trip was initiated now (based on the users known travel speed, and an estimate of the delay to be encountered at each signalized intersection along the way).

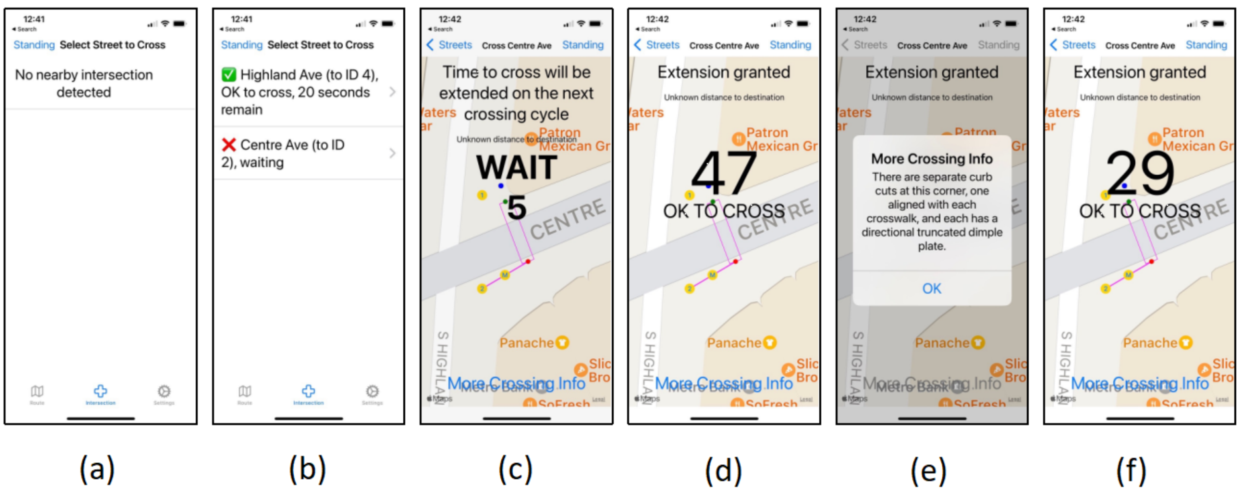

Figure 5: Crossing the street with ***PedPal+***

Once underway, the app is designed to remain in *route* mode as the user begins to execute the route. As the user progresses, the app is designed to track progress and provide navigation advice. These capabilities remain under development at the point of this report. When it is detected that the user is within range of the next signalized intersection, the app shifts to *intersection* mode and, once MAP and SPaT messages are processed, begins presenting crossing options (see Figure 5.a, 5.b). The user is subsequently informed of arrival at intersection corner (as determined through interaction with the UWB beacon at the corner). Once the user makes a crossing selection, the app communicates the required crossing time to the traffic signal system and is informed whether the crossing time extension has been granted. If the selection is for a future crossing phase, the app signals for the pedestrian to wait (Figure 5.c). As the time until the crossing phase nears, the app begins to count down the time remaining to prepare the user for the cross. When the crossing phase commences, the app announces that it is OK to cross (Figure 5.d) and begins to count down the time remaining in the phase. When the user is detected at the destination corner (again via interaction with the UWB beacon positioned at that corner), the app informs the user that the crossing is complete (Figure 5.3), and a new set of crossing options are presented. When the app detects that the user is moving away from the intersection, the interface reverts to *route* mode for navigation support.

2.5 Towards a Wearable Version

One use issue identified in the user pilot test of the original PedPal app stemmed from the fact that a smartphone app is difficult to take advantage of for pedestrians with certain types of disabilities. Wheelchair users, for example, have their hands occupied and would require some sort of mounting device to exploit PedPal (or ***PedPal+***) assistance. Likewise, older pedestrians that travel by means of a walker or a rollator do not have their hands free to interact effectively with the app.

To address this shortcoming, an initial port of the PedPal app to the Apple watch was undertaken. Given the reduction in screen size available on the watch, a wearable app design was produced that incorporates connectivity to the original smartphone app as a means of limiting the functionality provided by the wearable app to core support for navigating complete trips in urban environments. Specifically, *settings* mode is accessible only on the smartphone app, and the offline actions taken to configure and personalize the app for its user are retrieved by the wearable app when needed during operation. An initial version of the wearable app that provides basic street crossing functionality, depicted below in Figure 6, remains under development as of the writing of this report.

Figure 6: An initial wearable version of ***PedPal+***

2.6 ***PedPal+*** in Operation

To debug and validate the extended capabilities developed over the course of the project, a series of live tests were conducted at the intersection of Centre Avenue and Highland Avenue in the East Liberty neighborhood of Pittsburgh, culminating in execution of the `complete trip' demonstration scenario depicted in Figure 7. At a high level, the demonstration trip starts at a pedestrian's "home" location on Highland Avenue, proceeds north across Centre Avenue, then crosses Highland Avenue, and finally proceeds north to the bus stop on the corner of Penn and Highland Avenues, the next signalized intersection. Both intersections are part of the City's Surtrac traffic signal system deployment, and each corner of the intersection of Centre Avenue and Highland Avenue were additionally instrumented with iPhones serving as UWB beacons.

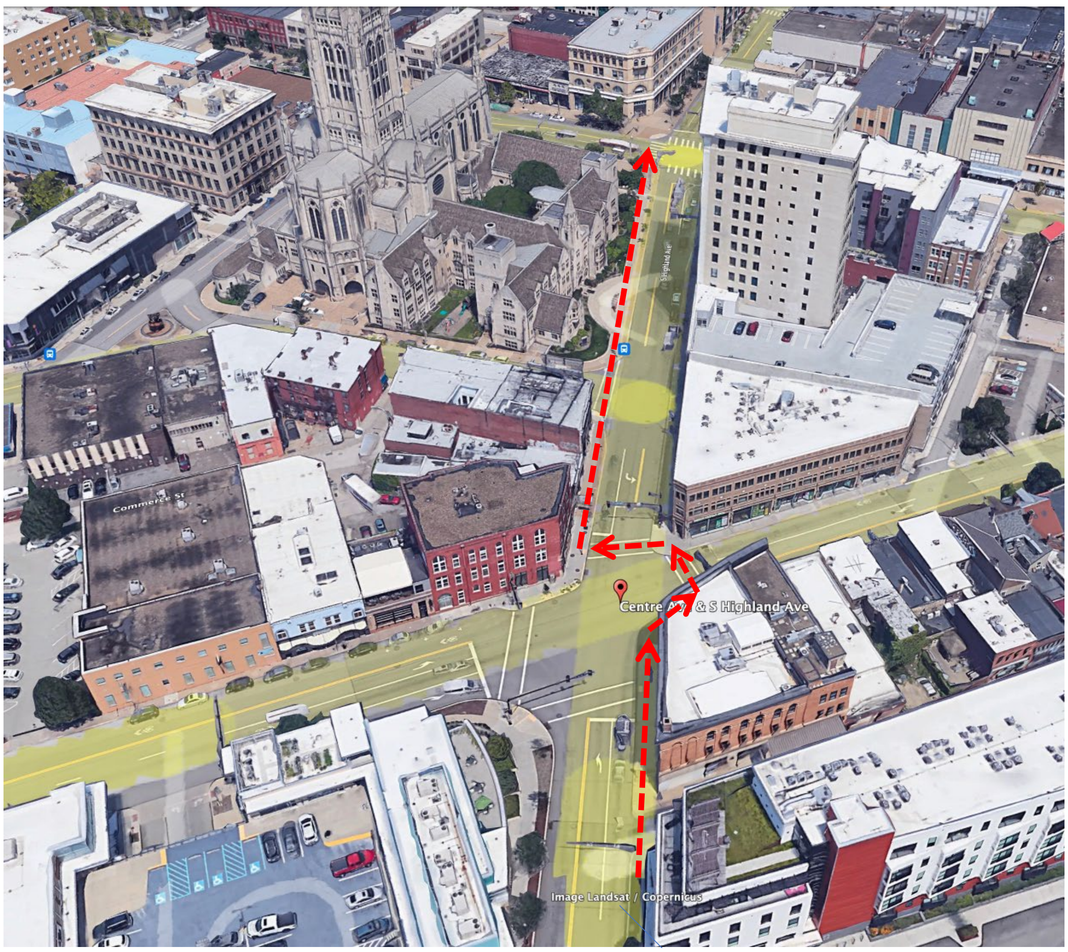

Figure 7: Demonstration scenario.

In more detail, the demonstration scenario was designed to incorporate the following progression of events:

1. The pedestrian first requests a route from his current location to the eastbound bus stop at Penn Avenue and Highland Avenue from the app. The app responds with the route shown in Figure 7 and an estimated travel time of x minutes (using its knowledge of how fast the user travels, how many signals need to be crossed, etc.).[2] Once the route has been generated, the user confirms that there is sufficient time to get to the bus top before the next bus traveling east is expected to arrive and proceeds to start along the route.
2. As the pedestrian approaches the intersection of Highland and Centre, the app recognizes the intersection corner (the southeast one in this case) through interaction with the Ultra-Wide Band (UWB) beacon positioned at the corner and presents the user with possible crossing options. Through interaction with the app (as indicated previously in Figure 5), the user crosses Centre Avenue moving north.
3. As the pedestrian begins to cross Centre, the app broadcasts the user's presence in the crosswalk to any approaching vehicles with connected vehicle capability. In the demonstration scenario, an ARGO AI vehicle synchronized to be traveling north on Highland from Shadyside and turning east onto Centre Avenue at the same time as the cross, receives the message that the user is in the crosswalk and slows down to allow the user to cross first.
4. Once across to the northeast Corner, the pedestrian is presented with new street crossing options and prepares to cross Highland Avenue to get to the NW corner, again using the app screen progression in Figure 5 (or its voiceover equivalent in the case of a vision impaired individual).
5. Although not yet operational, an anticipated short-term extension of the demonstration scenario will involve the user veering toward the intersection as he proceeds to cross Highland Avenue and stepping outside of the crosswalk, at which time the app sends a haptic signal to alert the user (e.g., single vibration if outside to the left, double vibration if outside to the right), and initiate course correction.
6. Once across the street, the user turns right and heads up Highland Avenue toward Penn Avenue and the bus stop.

For the final demonstration conducted at the end of the project, a blind individual was recruited to play the role of the pedestrian. A video recording of this demonstration can be found at *<link to be added>.*

[2] In actuality, the pathVu accessible routing procedure did not produce the route depicted in Figure 7. Instead, it judged the intersection to be too dangerous for the pedestrian to cross, and returned the route shown previously in Figure 4.c, which skirts the intersection and crosses at the next intersection traveling east bound on Centre Avenue, then follows a side street north to Penn Avenue and then heads west on Penn Avenue to the destination bus stop. To allow us to test and demonstrate at this more challenging intersection, we introduced an option to preload a route, and used it to preload the route depicted in Figure 7 for testing and the actual demonstration.

2.7 Next Steps

Our immediate next steps are to refine and harden the ***PedPal+*** technology and then pursue general deployment of it. One important prerequisite to deployment is maturation of the UWB beacon technology used for corner identification and pedestrian tracking. Substitution of newly available Estimote UWB beacons (or a suitable alternative) for the iPhone12 Minis used to validate the UWB technology solution to the localization problem will require some amount of reengineering. Specifically, use of iPhone Minis as UWB beacons allowed us to utilize cell phone connections to the ***PedPal+*** cloud server to simplify the process of corner detection, which is not possible with independent UWB beacon technologies that do not have cellular communication capabilities. One short-term objective will be to redirect this communication to the cloud server through the smartphone app itself.

A second short-term focus will be demonstrating the mechanism we have developed for detecting pedestrian movement outside of the crosswalk during crossing in the field. We believe this is basically a matter of developing a `nearby interaction' implementation of the pedestrian localization and reorientation algorithm summarized in Section 2.2 and should be straightforward once the transition to independent UWB beacon technology is complete. A final short-term focus will be development of a fully functional real-time route navigation system. The current app provides a rudimentary capability to track user progress but gives no pro-active auditory guidance to the user in navigating the route.

Several avenues present themselves for subsequent deployment of the ***PedPal+*** complete trip technology. The most likely paths will be through our deployment partners, and to promote this eventuality, all technology results produced under this effort have been designated as open source. This will allow PathVu to directly exploit the mobile app's integration with their sidewalk mapping database, to further develop its integrated path planning and safe intersection crossing functionality, and to integrate the ***PedPal+*** app into its current product offerings. Complementary progress with techniques for obstacle detection (see Section 3.2 below) will also give PathVu the option to further automate the collection of sidewalk data through adoption of robotic wheelchair sensing and obstacle detection technologies.

Similarly, Rapid Flow Technologies is incentivized to incorporate ***PedPal+***'s extended component capabilities for safe intersection crossing. Rapid Flow's basic approach to deployment of the original *PedPal* technology (which we expect to be the same for ***PedPal+***) is to provide it free of charge to municipalities who purchase the Surtrac traffic signal control system, on the assumption that municipalities could then offer the capability to their local disability community as a gesture of goodwill.

A final prospective path to deployment that we are pursuing is the concept of *PedPal-lite*, a version of the smartphone app that does not require Surtrac traffic signal control and instead interfaces with a conventional traffic controller. In this case, the ability to dynamically extend the crossing phase if the pedestrian is moving slower than expected is given up, but the resulting smartphone app becomes much more broadly deployable and most other capabilities of ***PedPal+*** remain available. We are initially targeting an NCTIP controller interface to maximize US deployment opportunities, and we envision a deployment strategy that combines a policy of free release of the app to any interested municipality with an optional service agreement for maintenance and upgrades.

3. Autonomous Wheelchair Navigation using Learned Traversability Model

The second major thrust of the `Complete Trip' project has focused on the development of enabling technology for providing motorized wheelchair users of the ***PedPal+*** app with the additional complementary capability to invoke autonomous driver-assist for traversing difficult route segments such as intersection curb cuts, sidewalks under repair, and potholes. To this end, a full-stack autonomous wheelchair (depicted in Figures 10 and 13 below) was developed as a research platform for investigating solutions to two core technical problems: outdoor navigation and dynamic obstacle detection. In this section, we summarize work performed to address the first problem. In Section 4, we describe technology that was developed to solve the second problem.

Autonomous wheelchair navigation involves the notion of shared autonomy where wheelchair users expect vehicles to provide safe and comfortable rides while following users' high-level navigation plans. To find such a path, vehicles negotiate with different terrains and assess their traversal difficulty. Most prior work models surroundings either through geometric representations or semantic classifications, which do not reflect perceived motion intensity and ride comfort in downstream navigation tasks. We instead focus on ride comfort explicitly in traversability analysis using proprioceptive sensing. In our approach, known here as *Ride Comfort-Aware Visual Navigation (RCA)*, a self-supervised learning framework has been developed to predict traversability cost map from first person-view images by leveraging vehicle states as training signals [Xinjie et. al 2022, Xinjie 2022]. Our approach estimates how the vehicle would "feel" if traversing over based on terrain appearances. We demonstrate that our navigation system provides human-preferred ride comfort through both experiments carried out on a prototype automated wheelchair and a human evaluation study. In the following subsections we elaborate elements of the technical approach taken and summarize the technical results obtained.

3.1 Problem Formulation

The target problem is that of predicting the navigation costs using first-person view, monocular camera images. Let ⟩t ∈ **I** denote an image input at time step t where **I** denotes the projection space of the monocular camera images. We aim to find a function that maps this projection space to a 2D cost map, denoted by $\Phi(I) \rightarrow R^2$. Based on the intuition that traversability directly affects the vehicle state, we introduce vehicle state as an internal representation of cost and formulate this problem as one of finding two mapping functions. Let s_t denote the vehicle state at time t, consisting of the 3D robot pose including position, orientation, angular velocity, and linear acceleration. The first subproblem is to find a mapping function $\Phi_{I \rightarrow S}(I) \rightarrow S$ that can predict vehicle states from input images, minimizing the error between the predicted and the true states. The second mapping function, $\Phi_{S \rightarrow R}(S) \rightarrow R^2$, maps vehicle states to traversability costs, quantifying the traversability based on vehicle states. Using this formulation, our self-supervised approach automatically generates training data in the form of an image and the corresponding traversability cost map. This training data can then be used to learn the mapping function from image to cost $\Phi(I) \rightarrow R^2$.

3.2 First-Person View Visual Navigation

Figure 8 shows a high-level view of the first-person view navigation system where local planning determines the smoothness of a navigation path. As illustrated here, the terrain traversability analysis module uses first-person view camera images to generate a traversability cost map for the local planner to decide the actual navigation path. We compare our proposed RCA approach with two baseline approaches, namely 3D-based and Semantic-based approaches.

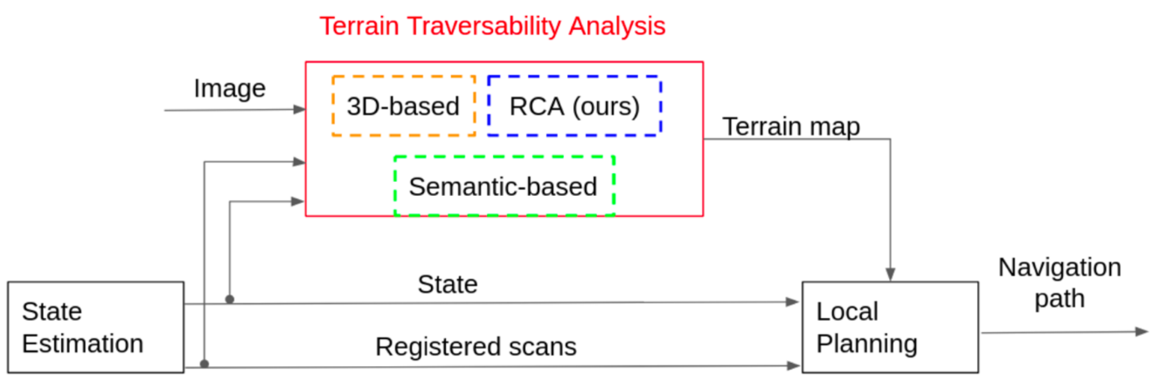

Figure 8: A navigation system with the terrain analysis module. Our contributions center on the terrain traversability analysis module that provides guidance for local planning.

3.3 Ride Comfort Aware (RCA) traversability analysis approach

Figure 9 provides an overview of the RCA approach. During training, the vehicle experiences physical vibrations by traversing various types of terrains. Using recorded vehicle states and first-person view images, the system learns to estimate traversability cost from images. Instead of manual labeling, we first use unsupervised learning to coarsely group vehicle states into clusters based on vehicle-terrain interactions. Next, with those clusters, we define a continuous cost function to reflect the traversability based on vehicle dynamics. Finally, we train a prediction model to associate terrain images directly to the traversability costs. During deployment, the learned model estimates a terrain cost map to support autonomous navigation. Each of these three steps are considered in more detail in the subsections below.

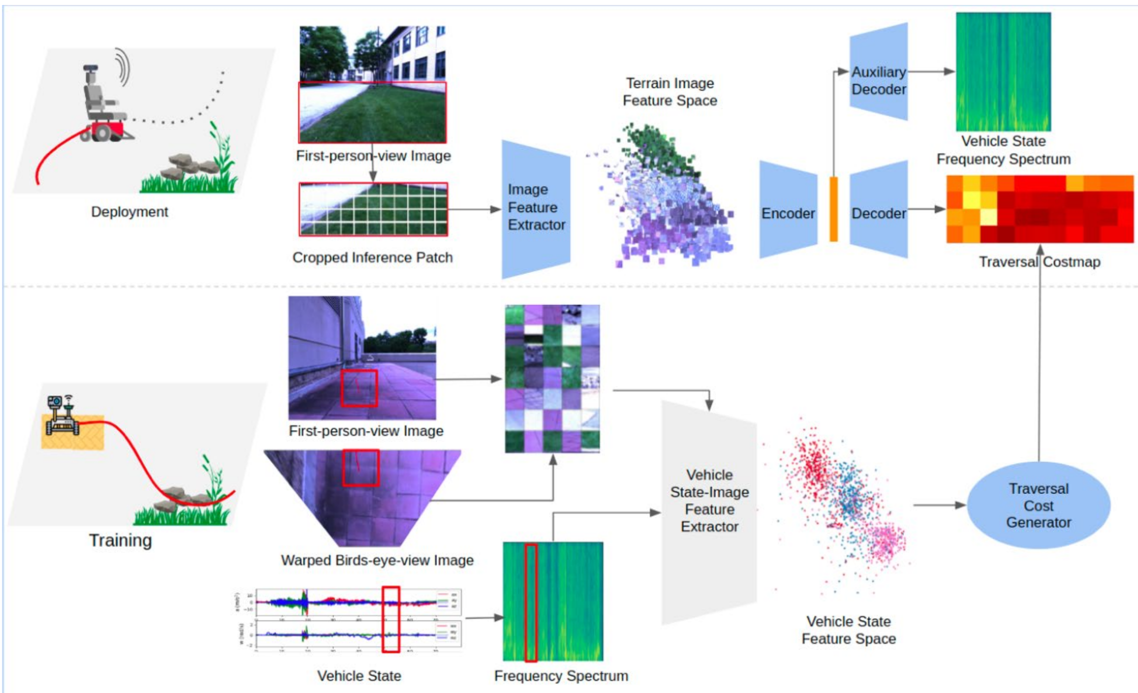

Figure 9: An overview of the Ride Comfort Aware (RCA) Visual Navigation Approach

3.3.1 Estimating Traversability Cost from Vehicle States

Within the RCA framework, vehicle states are utilized only during training as intermediate learning signals and as indicated above, the first step is to use them to estimate traversability costs. By turning these learning signals into numerical costs, the principal learning algorithm to be applied subsequently can use these predicted costs, along with the images, for self-supervision, rather than utilizing vehicle states directly. Estimation of traversability costs from vehicle states is itself accomplished in two steps. First, an unsupervised learning method is used to generate coarse labels for vehicle state clusters. Second, the learned clusters are used to estimate a continuous cost function given a vehicle state input.

<u>Step 1: Unsupervised Learning for Classifying Vehicle States</u>. After input data preprocessing, we have image patches indicating traversal regions and frequency spectra describing experienced motions. As opposed to previous self-supervised approaches to vibration-based classification which have handcrafted features, we exploit the complementary information shared by the visual and vehicle state domains to extract latent features. To learn this association, we generally follow an existing unsupervised feature learning framework for classifying acoustic data [Zurn 2021]. In essence, this framework brings visually similar samples closer and visually distinct samples further away in the target space. We customize this framework for interpretation of riding comfort as follows. First, we extract image features using the Deep Encoding Pooling (DEP) network [Hue 2018], which is specialized for ground terrain recognition. Next, to ensure that negative samples are selected from different ground truth classes, we leverage semantic classes clustered in the visual feature space. Such prior knowledge serves as a reference for computing Euclidean distance and selecting negative samples. Based on the Principal Component Analysis (PCA) projection of the feature space, we perform k-Means clustering to obtain coarse labels for terrain classes.

<u>Step 2: Traversal Cost Generation</u>. Wheelchair passengers are exposed to greater physiological risks and psychological barriers with drastic changes in the vehicle motions. Thus, we hypothesize that larger and more frequent movements along any of the three dimensions of *roll, pitch*, and *Z* should be avoided and assigned higher cost. The amplitude spectra show amounts of motion variations at different frequencies, and vehicle state clusters obtained in the previous step serve as a prior to show distinct vehicle dynamics. The traversal cost function is considered as a weighted average of roll, pitch, and Z, measuring similarity to the average amplitude spectrum of its vehicle state class. Different vehicle state classes are also magnified with a hyperparameter to offset from each other. Let $A_i^{d,k} \in R^m$ denote the amplitude spectrum along dimension d from the vehicle state class k for the sample i. N^k describes the number of samples within the vehicle state class k. A traversal cost for sample i from vehicle state class k is,

$$T_i^k = \omega_k \cdot \sum_{d=1}^{3} \mathbf{M^{d,k}} \cdot \mathbf{A_i^{d,k}}$$

where ω_k is a weight parameter of vehicle state class k, and where $M^{d,k}$ is the mean amplitude spectrum along dimension d from the vehicle state class k, i.e.,

$$\mathbf{M^{d,k}} = \frac{1}{N^k}\sum_{i=1}^{N^k}\mathbf{A_i^{d,k}}$$

3.3.2 Self-supervised Traversal Cost Prediction

With the traversal cost function, we associate terrain visual appearances with motion measurements. We implement an encoder-decoder network, to estimate amplitude spectra of vehicle states and to predict traversal costs. To align the visual feature space with the vehicle state feature space, terrain patches are first extracted by the same feature extractor. Then terrain visual features are further encoded. During training, two decoders trigger back propagation at different epochs. In early epochs, the encoded space is translated toward vehicle states and only weights from the auxiliary vehicle state decoder get updated. Then later, the traversal cost decoder joins and learns from the partially translated feature space. We use L2 loss and smooth-L1 loss for regressing amplitude spectrum and traversal cost, respectively.

$$\mathcal{L}_2 = \left\| g\left(f\left(\mathbf{I}_i\right)\right) - \mathbf{A}_i \right\|^2$$

Define x = $h(f(\mathrm{I}_i)) - T_i$ as the difference between the traversal cost estimation and the computed traversal cost,

$$smooth_{\mathcal{L}_1} = \begin{cases} 0.5x^2 & \text{if } |x| < 1 \\ |x| - 0.5 & \text{otherwise} \end{cases}$$

The regression loss is written as,

$$\mathcal{L} = \beta\mathcal{L}_2 + (1-\beta)smooth_{\mathcal{L}_1}$$

Given a camera image, it infers the traversal cost for each cropped patch and aggregates them as 2D cost maps.

3.4 Results

3.4.1 Experiment Design

<u>Dataset</u>: We collected a training dataset in an urban area using a tracked vehicle (Fig. 5.1a) that competently covered a variety of challenging terrains without posing risks to human operators. The camera captures frontal images at 5Hz with 1280 × 1024 resolution. The state estimation module provides motion measurements at 200 Hz. Our

dataset includes 2 hours of operated driving at 1 m s−1 including variations of soft and hard surfaces with different elevations to ensure a broad spectrum of vehicle motion profiles.

Evaluation: To evaluate the quality of RCA, we compared the performance with two baselines: a 3D-based approach [Cao 2021] and a semantic classification-based approach [Wu 2021]. To assess the performance of three terrain analysis approaches in actual navigation tasks, we conducted robot experiments with a wheelchair-based vehicle navigating various terrain conditions. To investigate the performance according to various factors determining ride comfort and to evaluate perceived motion profiles, we designed and conducted a human evaluation study.

Figure 10: The wheelchair robot platform used in our experiments is equipped with a 4.1 GHz i7 computer and a NVIDIA GTX 1660Ti GPU card.

3.4.2 Robotic Wheelchair Performance

The three navigation approaches identified above (RCA, 3D-base, semantic classification) were tested in five scenes including asphalt roads, grass, curbs, tactile pavings, gravels (see Figure 11, top row). The wheelchair starts at the same initial pose and attempts to reach the desired goal points (designated by a red cross). The bottom row of Figure 11 shows recorded wheelchair trajectories in a top-down view (our RCA in blue; semantic-classification in red; and 3D-based in green). In Figure 11.1, RCA detours away from the curb and uneven surfaces until it reaches the ramp. In Figure 11.2 RCA takes a side way while avoiding the turbulence brought by tactile pavings. In Figures 11.3 and 11.4, RCA foresees the effect from sinking in the soft terrain and avoids the vegetation although going through it is the shortest path. In Figure 11.5, RCA keeps the wheelchair away from the vegetation and gravels.

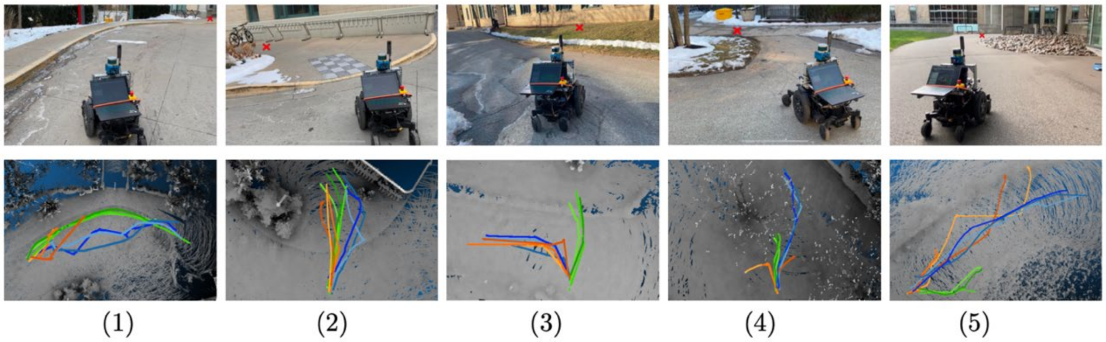

Figure 11: Robotic wheelchair performance in the field.

3.4.3 Perceived Motion Intensity versus Human Evaluation

We designed an Amazon Mechanical Turk (AMT) study to assess human evaluation of the performances on five dimensions of ride comfort. During a session, workers first watched three videos showing different robot behaviors in the same scenario. The order of these three videos is randomized. The participants were then asked to rank the three videos from the most preferred (1) to the least preferred (3) according to the following five criteria:

- Stability: A ride is considered stable when noticing few frame-to-frame jitters and gentle orientation changes.
- Path normality: A ride is considered normal if the path matches with a human's expectation.
- Safety: A ride is considered as safe when observing continuously smooth motions, predictable paths, and avoiding obstacles at proper distance.
- Trustworthiness: A ride is considered as trustworthy if a passenger could rely on the wheelchair to complete trips independently.

A ride is most preferred when you would like to have a same wheelchair in your community.

Table 1 below reports the Spearman Correlation scores between stability rankings and Perceived Motion Intensity (PMI) ranks [de Winkel 2020] along Z, Roll, Pitch, Yaw axes among all experiment runs with $p < 0.01$. It reveals strong positive correlations between human perceived stability scores and Perceived Motion Intensity computed from acceleration and jerk along Z, Roll, Pitch, Yaw axes. It further validates that our approach improves the perceived stability relatively by reducing motion intensity along Z, Roll, Pitch, Yaw axes.

	PMI Z	PMI Roll	PMI Pitch	PMI Yaw
Stability	0.69	0.70	0.61	0.80

Table 1: PMI versus Human Evaluation

Figure 12 presents a vote breakdown in percentage received by each approach among all participants. Each vertical bar consists of three segments representing the percentage of samples voting an approach as the first (bottom, darker color), the second (middle, moderate color), the third (top, lighter color). Our approach marks as the blue bar, while the semantic-based baseline is in red, and the 3D-based baseline is in green. Across five evaluation criteria over five scenes, our approach is most frequently ranked as the top. The results show no clear difference between semantic-based approach and 3D-based approach. The results also cast light on the positive correlation between objective factors (i.e., stability, safety) to subjective factors (i.e., trustworthy, preference).

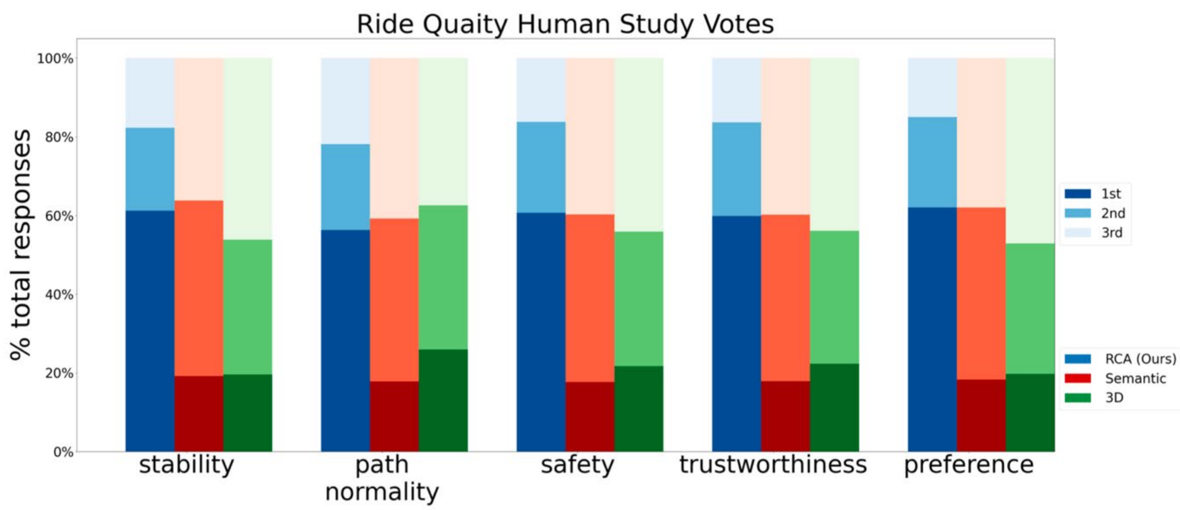

Figure 12: Percentage of samples ranking three approaches among five dimensions.

3.4.4 Summary

Based on human evaluation of 55 AMT workers, RCA is consistently ranked the highest when compared to the baseline approaches in terms of stability, path normality, safety, trustworthiness, and overall preference. We also recognize that RCA does not perform consistently well across all scenarios. Whereas RCA generally performs well by estimating the disturbing motions brought by uneven or soft terrains, the semantic-based approach provides human-interpretive estimation in the pixel space and the 3D-based approach provides robust and accurate geometric measurements of rigid objects. These findings lead to our future direction for unifying comfort-aware, semantic-based, and 3D-based approaches towards the goal of supporting safe and comfortable rides for the wheelchair users.

4. Dynamic Object Recognition for Real-Time Obstacle Avoidance

A second challenge associated with autonomous wheelchair navigation is that of detecting and avoiding obstacles, some of which (e.g., other pedestrians) may be in motion and rapidly approaching. One final research thrust with the `Complete Trip' project focused on the development and demonstration of a novel 3D adaptive safety sensor that enables dynamic obstacle detection and avoidance. This sensor technology, referred to as "*Programmable Light Curtains*", was invented by the imaging group at Carnegie Mellon University. It is shown in Figure 13 below as the blue box mounted on the full-stack autonomous wheelchair mentioned earlier as our research platform. Over the course of the project, novel algorithms based on this sensing technology were developed for wheel-chair localization and odometry, mapping of the stationary scene, dynamic obstacle identification and avoidance and finally path planning for reaching the destination. The methods and system were developed by a PhD and a MS student and became important parts of their theses. The autonomous navigation that resulted was demonstrated in cluttered environments like the university labs, corridors, and cafeterias. [3]

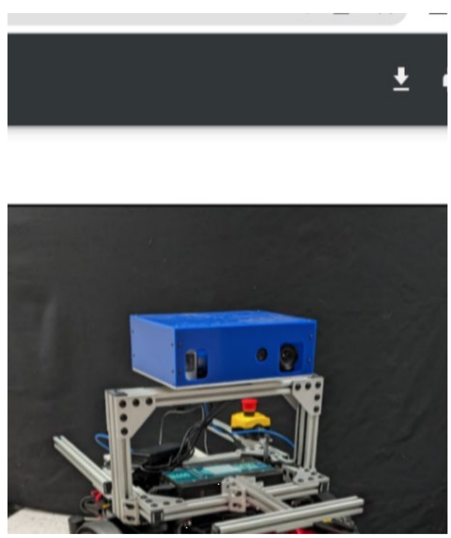

Figure 13: Full-stack autonomous wheelchair with Programmable Light Curtain sensor attached.

Most autonomous systems today rely on LIDARs for 3D perception. While these systems have been very successful, there are several important hurdles in broadly

[3] For further information about this technology, please see https://www.cs.cmu.edu/~ILIM/light_curtains

using them. First, LIDARs are too expensive compared to the cost of a wheelchair. Second, they provide only low-resolution and low-framerate 3D information and may miss small obstacles that may be important to account for while navigating in cluttered environments. Third, they require large computing resources for obstacle detection, tracking and avoidance. Our novel sensor (shown in more detail below in Figure 14a) provides a cost-effective and reliable safety solution for wheelchair navigation in busy environments. Programmable light curtains (PLC) provide high resolution and high frame-rate 3D information. They can see through dust, fog, or rain, and provide direct and efficient obstacle identification capabilities. LIDAR based systems on the other hand are expensive, are typically low resolution missing small objects and low-frame rate. This safety sensor is now being commercialized by the company Phlux Technologies, a CMU start up. Discovery research and fellowship for the CEO of the startup was partially sponsored by US DOT funding. The example application visualized in Figure 14b shows the use of the light curtains between a street and a sidewalk that detects pedestrians that are about to step onto the street. No additional computation is required here.

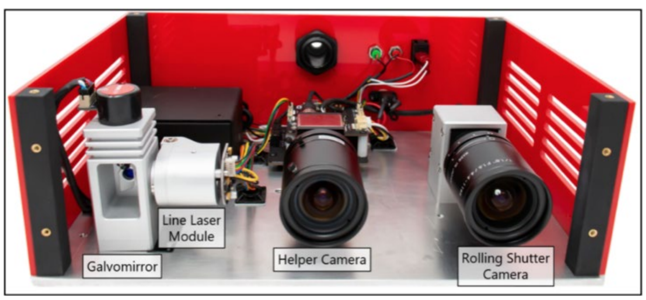

(a)

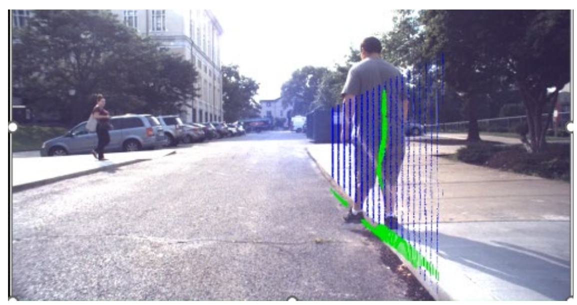

(b)

Figure 14: (a) Programmable Light Curtain (PLC) sensor; **(b)** PLC detection of pedestrian movement from the sidewalk into the street.

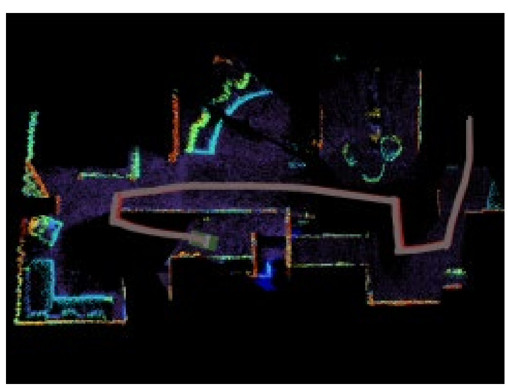

Figure 15: A sample navigation run through a cluttered indoor environment

The first stage in the autonomy stack is the localization and odometry computation of the wheelchair. We developed an approach that combines the visual features from the camera and the 3D information in the programmable light curtain (PLC). We demonstrated accuracy of localization to be within a few centimeters – rivaling the systems that use LIDARs. An example navigation run in a cluttered home environment is shown in Figure 15. The solid line depicts the path that the wheelchair took. Pictured in Figure 16 is a cluttered office corridor and the 3D map that is computed in real-time using our novel Simultaneous Localization and Mapping algorithm. The wheelchair successfully navigates in these environments.

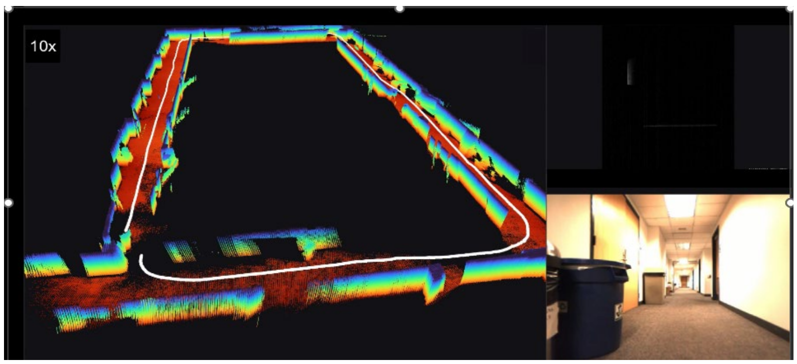

Figure 16: Navigation of the wheelchair in a narrow office corridor that is cluttered. Our system computes the map of the scene in 3D, localizes the wheelchair in the map and avoids any obstacles.

Once the map of the scene is computed, the system uses it to identify dynamic obstacles, like people or vehicles and plans for a path that avoids collisions. The result on the right shows the wheelchair navigating in the presence of people. The obstacle avoidance unit produces multiple safe paths (illustrated in Figure 17 below), and the planner identifies the most efficient path to reach the goal.

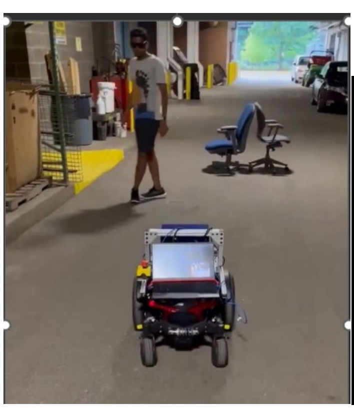

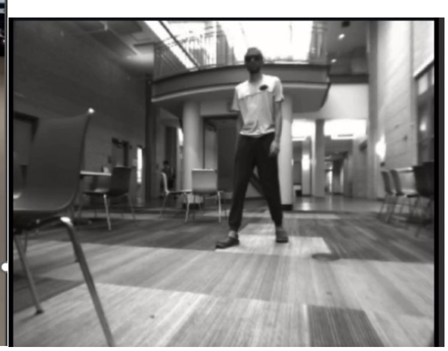

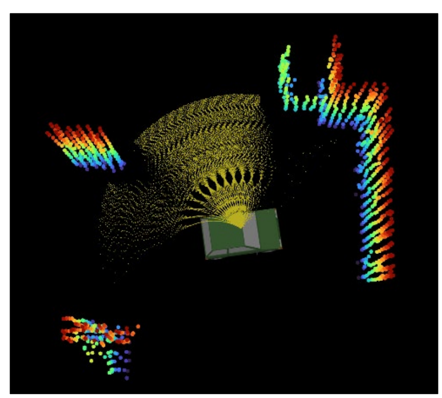

Figure 17: The wheelchair navigating around a person in a corridor and an office cafeteria. The path planner identifies obstacles in real-time and identifies the most efficient path around the person. The programmable light curtain achieves this with minimal computations as compared with traditional approaches that use either cameras or LIDARs*.*

5. Publications and Presentations

5.1 Publications

- Ancha, Siddhartha, "Active Robot Perception using Programmable Light Curtains", PhD thesis, The Robotics Institute, Carnegie Mellon, 2021.
- Ancha, S., Pathak, G., Held, D, Zhang, J., Narasimhan, S. G., "Active Velocity Estimation using Light Curtains via Self-Supervised Multi-Armed Bandits", Submitted for review to *RSS 2023*.
- Hata, R. Isukapati, I., Rubinstein, Z.R., and Smith, S.F., "Blind Pedestrian Localization and Reorientation at Urban Crosswalks via Ultra-Wide Band Beacons, *Proceedings 2022 CMU Robotics Institute Summer Scholars (RISS) Program*, August 2022.
- H-C Hu, G.J. Barlow, J. Zhou, and S.F. Smith, "CARIC: Connection-Based Scheduling for Real-Time Intersection Control", unpublished working paper, November 2022 (submitted to *ICAPS 2023*)
- Neiman, D., Z.B. Rubinstein, and S.F. Smith, "Dynamic Route Guidance in Vehicle Networks by Simulating Future Traffic Patterns", unpublished working paper, November 2022 (submitted to *ICAPS 2023*)
- Pathak, Gaurav, "Programmable light curtains for Safety Envelopes, SLAM and Navigation", MS Thesis, The Robotics Institute, Carnegie Mellon, 2021.
- Smith, S.F. "Routecast: Integrating Connected Vehicle Technology with Adaptive Traffic Signal Control to Revolutionize Urban Mobility", Rapid Flow Technologies White Paper, February 2021.
- Xinjie Yao, Ji Zhang, Jean Oh. "RCA: Ride Comfort-Aware Visual Navigation via Self-Supervised Learning", *35th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'22)*, Kyoto, 2022.
- Xinjie Yao, "RCA: Ride Comfort-Aware First-Person Navigation via Self-Supervised Learning" MS Thesis, The Robotics Institute, Carnegie Mellon University, CMU-RI-TR-22-15, 2022.

5.2 Presentations

- Smith, S.F., "Connecting Pedestrians with Disabilities to Traffic Signal Control for Safe Intersection Crossing and Enhanced Mobility, Transportation Research

Board’s Traffic Signal Control Committee Annual Meeting, Wash DC, January 9, 2022.

- “Smith, S.F. “Smart Transportation Infrastructure”, First University Transportation Center Video Webinar, US Department of Transportation, February 24, 2022 (https://youtu.be/L8p4pGYKrqw)
- Smith, S.F., Making Smart Signals Smarter and Safer through Connectivity with Travelers”, Transportation Research Board’s Traffic Signal Control Committee Mid-Year Meeting, Beckman Center, Irvine CA, July 6, 2022.
- Smith, S.F., “Making Smart Signals Smarter and Safer through Connectivity with Travelers”, Invited talk, 2022 *Intelligent Transportation Systems World Congress*, Los Angeles, CA, September 18, 2022.

References

[Cao 2021] Chao Cao, Hongbiao Zhu, Fan Yang, Yukun Xia, Howie Choset, Jean Oh, and Ji Zhang. Autonomous exploration development environment and the planning algorithms. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021

[de Winkel 2020] Ksander N. de Winkel, Florian Soyka, and Heinrich H. Bu ̈lthoff. The role of acceleration and jerk in perception of above-threshold surge motion. March 2020. doi: 10.1007/s00221-020-05745-7.

[Hata et.al 2022] Hata, R. Isukapati, I., Rubinstein, Z.R., and Smith, S.F., "Blind Pedestrian Localization and Reorientation at Urban Crosswalks via Ultra-Wide Band Beacons, *Proceedings 2022 CMU Robotics Institute Summer Scholars (RISS) Program*, August 2022.

[Hawkes 2016] Hawkes, A., "Traffic Control with Connected Vehicle Routes in Surtrac", MS Thesis, The Robotics Institute, Carnegie Mellon University, 2016.

[Hue 2018] Jia Xue, Hang Zhang, and Kristin J. Dana. Deep texture manifold for ground terrain recognition. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2018.

[MCMM 2020] THE COMPLETE TRIP: Helping Customers Make a Seamless Journey, National Center for Mobility Management, 2020.

[PathVu 2022] http://www.pathvu.com

[Smith et.al 2013] Smith, S.F. G.J. Barlow, X-F Xie, and Z.B. Rubinstein, "Smart Urban Signal Networks: Initial Application of the SURTRAC Adaptive Traffic Signal Control System", *Proceedings 23rd International Conference on Automated Planning and Scheduling,* Rome, Italy, June 2013.

[Smith et.al 2019] Smith, S.F., Rubinstein, Z.B, Marusek, J., Dias, B, and Radewick, H., "Connecting Pedestrians with Disabilities to Adaptive Signal Control for Safe Intersection Crossing and Enhanced Mobility: Final Report", Technical Report FHWA-JPO-19-754, September 2019.

[Smith 2020] Smith, S.F., "Smart Infrastructure for Future Urban Mobility, AI Magazine, 41(1) Spring 2020.

[Smith 2021] Smith, S.F. "Routecast: Integrating Connected Vehicle Technology with Adaptive Traffic Signal Control to Revolutionize Urban Mobility", Rapid Flow Technologies White Paper, February 2021. https://f.hubspotusercontent20.net/hubfs/4469970/PDF%20Folder/Routecast/RapidFlow-Routecast-WhitePaper-March2021.pdf?utm_medium=email&_hsmi=113935456&_hsenc=p2ANqtz-8mrkbFk7gJA04jTNjME72JlCZt_zKLF_ZsDAN63AyZPrcJ0GRerNfXr-BHz4XT2ZVmtlaX0dJ9_cpBRHKamQ2x6KBd

[Wu 2021] Zhanxin Wu, Xinjie Yao, and Jean Oh. Semantic segmentation in complex scenes for robotics navigation. In *Carnegie Mellon Robotics Institute Summer Scholars Working Papers Journal*, 2021.

[Xie et. al 2012] Xie, X-F, S.F. Smith, L Lu, and G.J. Barlow, “Schedule-Driven Intersection Control”, *Transportation Research Part C: Emerging Technologies*, 24: 168-189, October 2012.

[Xinjie et. al 2022] Xinjie Yao, Ji Zhang, Jean Oh. “RCA: Ride Comfort-Aware Visual Navigation via Self-Supervised Learning”, *35th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’22)*, Kyoto, 2022.

[Xinjie 2022] Xinjie Yao, "RCA: Ride Comfort-Aware First-Person Navigation via Self-Supervised Learning" MS Thesis, The Robotics Institute, Carnegie Mellon University, CMU-RI-TR-22-15, 2022.

[Zurn 2021] Jannik Zurn, Wolfram Burgard, and Abhinav Valada. Self-supervised visual terrain classification from unsupervised acoustic feature learning. *IEEE Transactions on Robotics*, 37:466–481, 2021.